



US 20200357142A1

(19) **United States**

(12) **Patent Application Publication**
AYDIN et al.

(10) **Pub. No.: US 2020/0357142 A1**

(43) **Pub. Date: Nov. 12, 2020**

(54) **LEARNING-BASED SAMPLING FOR IMAGE
MATTING**

(71) Applicant: **Disney Enterprises, Inc.**, Burbank, CA
(US)

(72) Inventors: **Tunc Ozan AYDIN**, Zurich (CH);
Ahmet Cengiz ÖZTIRELI, Zurich
(CH); **Jingwei TANG**, Zurich (CH);
Yagiz AKSOY, Zurich (CH)

(21) Appl. No.: **16/408,199**

(22) Filed: **May 9, 2019**

Publication Classification

(51) **Int. Cl.**
G06T 7/90 (2006.01)
G06N 3/08 (2006.01)
G06N 20/00 (2006.01)
G06T 11/60 (2006.01)
G06T 11/00 (2006.01)
G06T 7/194 (2006.01)
G06T 3/40 (2006.01)


(52) **U.S. Cl.**

CPC **G06T 7/90** (2017.01); **G06N 3/08**
(2013.01); **G06N 20/00** (2019.01); **G06T 11/60**
(2013.01); **G06T 2207/20084** (2013.01); **G06T**
7/194 (2017.01); **G06T 3/40** (2013.01); **G06T**
2207/10024 (2013.01); **G06T 2207/20081**
(2013.01); **G06T 11/001** (2013.01)


(57)

ABSTRACT


Techniques are disclosed for image matting. In particular, embodiments decompose the matting problem of estimating foreground opacity into the targeted subproblems of estimating a background using a first trained neural network, estimating a foreground using a second neural network and the estimated background as one of the inputs into the second neural network, and estimating an alpha matte using a third neural network and the estimated background and foreground as two of the inputs into the third neural network. Such a decomposition is in contrast to traditional sampling-based matting approaches that estimated foreground and background color pairs together directly for each pixel. By decomposing the matting problem into subproblems that are easier for a neural network to learn compared to traditional data-driven techniques for image matting, embodiments disclosed herein can produce better opacity estimates than such data-driven techniques as well as sampling-based and affinity-based matting approaches.



A



B



C

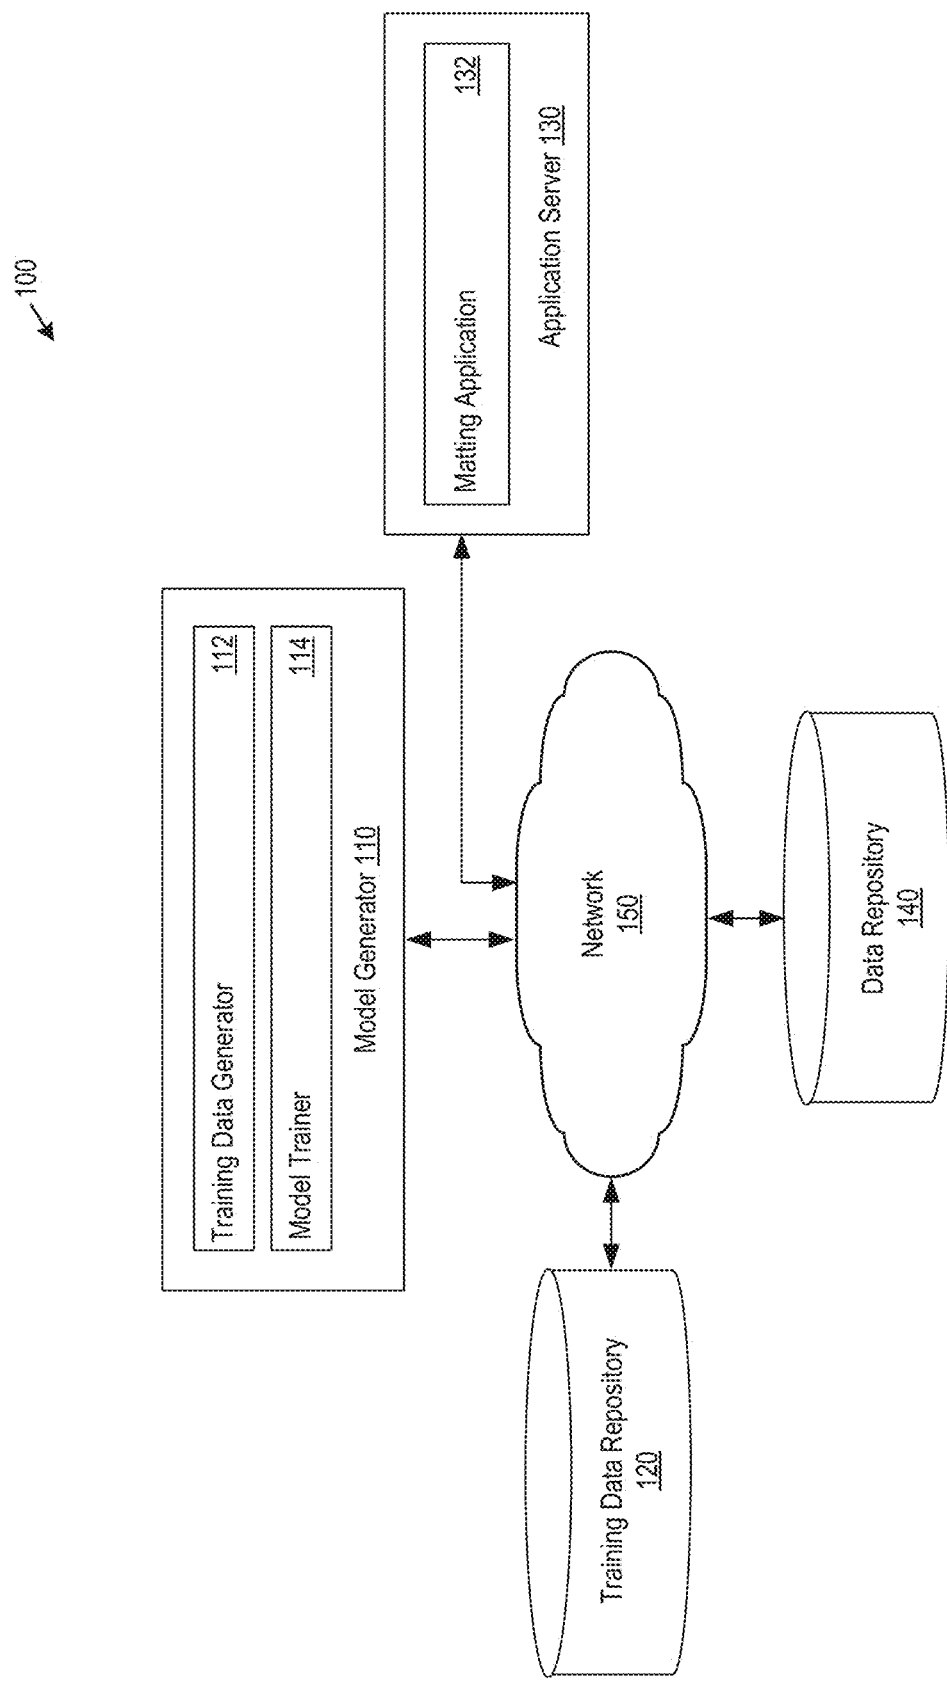


FIG. 1




FIG. 2






FIG. 3

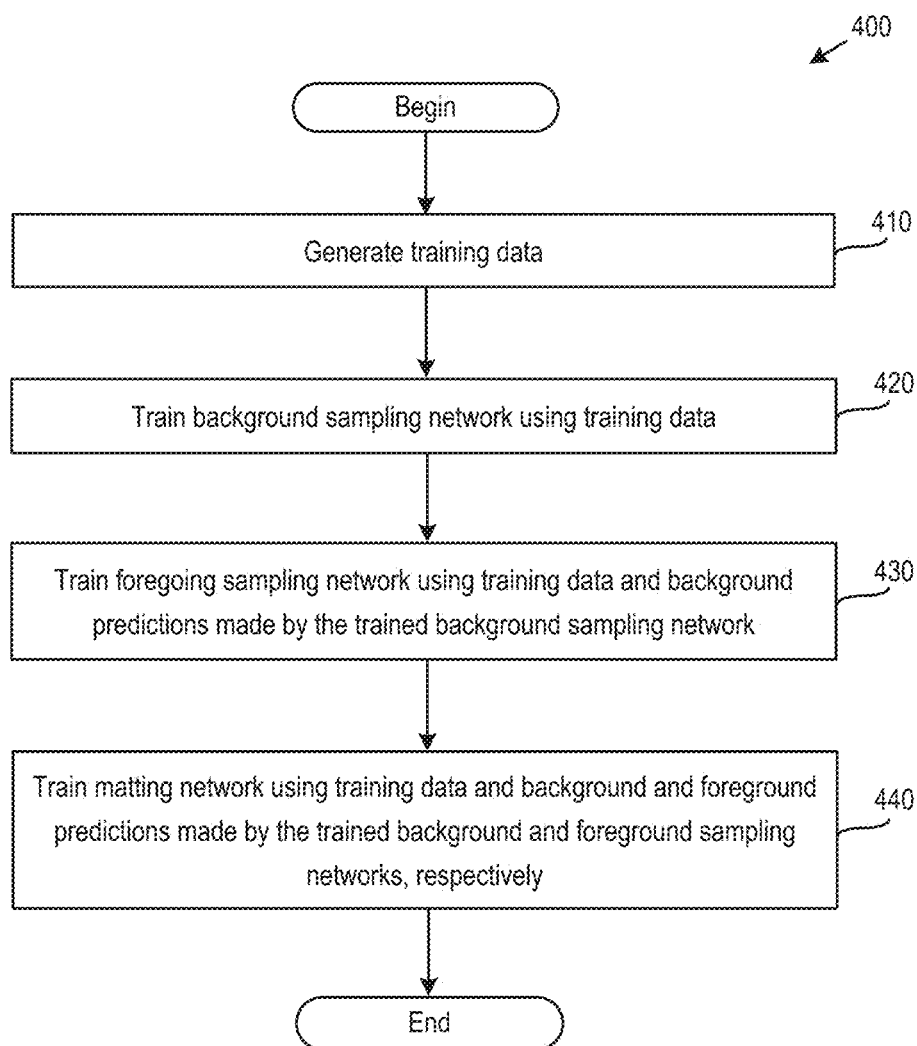
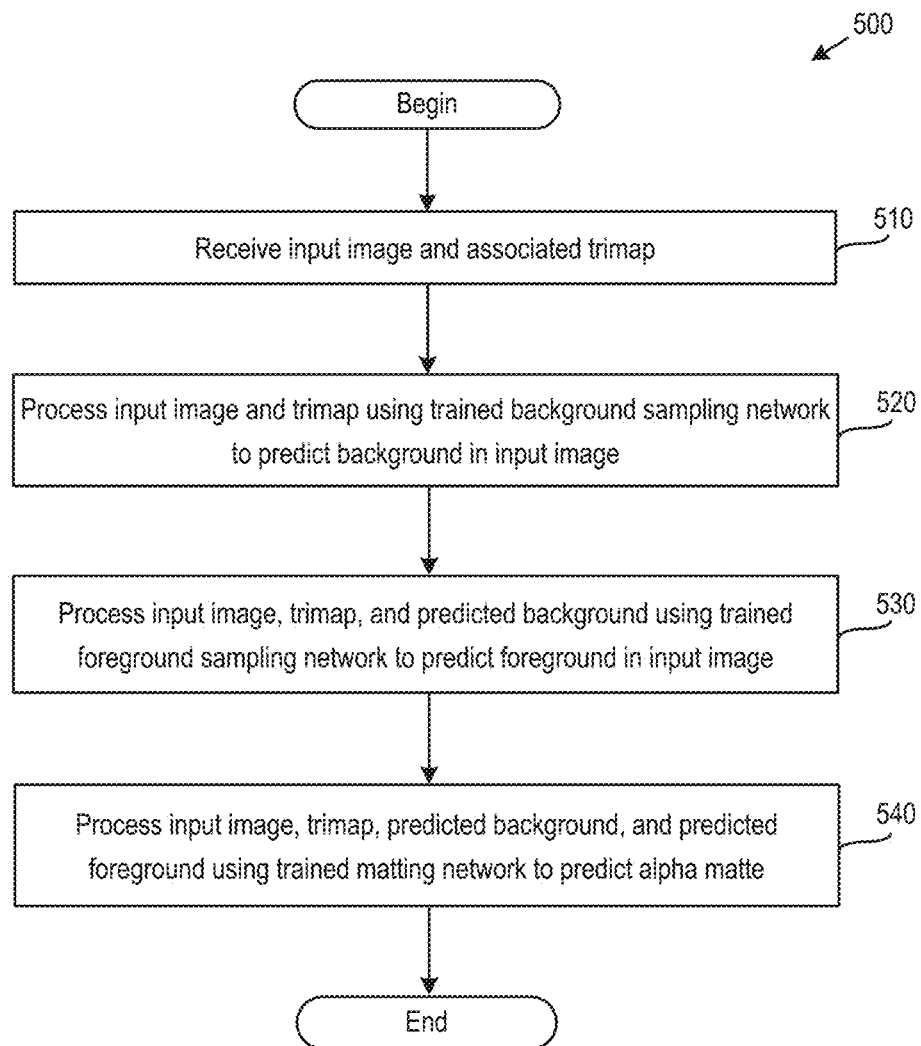


FIG. 4

**FIG. 5**

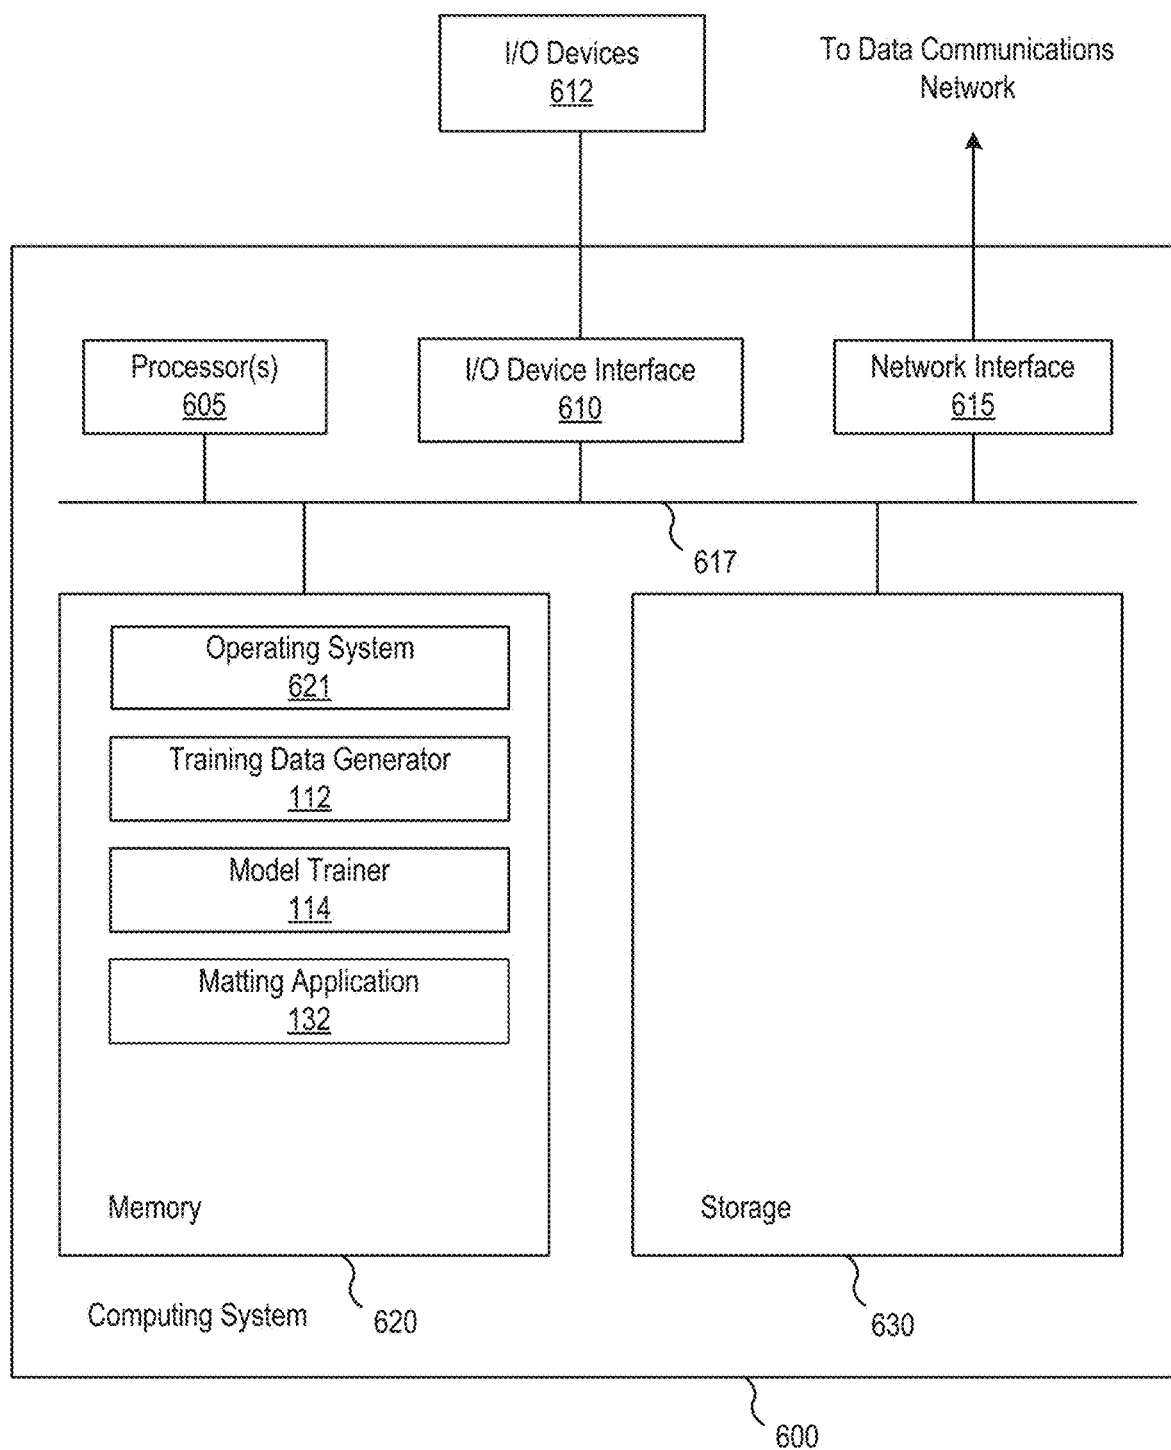


FIG. 6

LEARNING-BASED SAMPLING FOR IMAGE MATTING

COLOR DRAWINGS

[0001] The patent or application file contains at least one drawing executed in color. Copies of this patent or patent application publication with color drawing(s) will be provided by the Office upon request and payment of the necessary fee.

BACKGROUND

Field

[0002] This disclosure provides techniques for image matting.

Description of the Related Art

[0003] Image matting seeks to estimate the opacities of user-defined foreground objects in an (natural or synthetic) image, i.e., to estimate the soft transitions between foreground that is user-defined and a background, with the soft transitions defining the opacity of the foreground at each pixel. An alpha matte indicating the opacity of the foreground at each pixel (with, e.g., white indicating complete opacity, black indicating complete transparency, and shades of gray indicating partial transparency) may be obtained via image matting, and such an alpha matte is useful in many image and video editing workflows. For example, matting is a fundamental operation for various tasks during the post-production stage of feature films, such as compositing live-action and rendered elements together, and performing local color corrections.

[0004] One traditional approach to image matting is through sampling, in which color samples are gathered from known-opacity regions to predict foreground and background layer colors to use for alpha estimation. Such sampling-based matting typically involves selecting a color pair by using the color line assumption and other metrics from the spatial proximity of samples among others. However, an inherent shortcoming of the sampling-based matting approach is the lack of consideration for image structure and texture during the sample selection process. Further, current sampling-based matting approaches are limited to selecting only a single sample pair to represent a pixel's foreground and background colors.

[0005] Another traditional approach to image matting is affinity-based matting. Affinity-based matting techniques make use of pixel similarity metrics that rely on color similarity or spatial proximity to propagate alpha values from regions with known opacity. In particular, local affinity definitions, such as the matting affinity, may operate on a local patch around a pixel to determine the amount of information flow and propagate alpha values accordingly. However, affinity-based approaches can suffer from high computational complexity and memory issues.

[0006] As traditional approaches for image matting, such as sampling-based matting and affinity-based matting, tend to produce low quality results, current matting practices still require intensive manual labor.

SUMMARY

[0007] One embodiment disclosed herein provides a computer-implemented method for image matting. The method

generally includes processing a received image and associated trimap using, at least in part, a first machine learning model which outputs a predicted background. The method further includes processing the received image, the associated trimap, and the predicted background using, at least in part, a second machine learning model which outputs a predicted foreground. In addition, the method includes processing the received image, the associated trimap, the predicted background, and the predicted foreground using, at least in part, a third machine learning model which outputs an alpha matte.

[0008] Another embodiment disclosed herein provides a method for training machine learning models for image matting. The method generally includes training a first machine learning model using at least a portion of training data which includes images and corresponding trimaps. The method further includes training a second machine learning model using at least a portion of the training data and backgrounds predicted by the first machine learning model subsequent to the training of the first machine learning model. In addition, the method includes training a third machine learning model using at least a portion of the training data, backgrounds predicted by the first machine learning model, and foregrounds predicted by the second machine learning model subsequent to the training of the second machine learning model. The training of the third machine learning model attempts to minimize a loss function which includes a loss defined over alpha gradients.

[0009] Other embodiments include, without limitation, computer-readable media that include instructions that enable processing unit(s) to implement one or more embodiments of the above methods, as well as systems configured to implement one or more aspects of the above methods.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] So that the manner in which the above recited aspects are attained and can be understood in detail, a more particular description of embodiments of the invention, briefly summarized above, may be had by reference to the appended drawings.

[0011] It is to be noted, however, that the appended drawings illustrate only typical embodiments of this invention and are therefore not to be considered limiting of its scope, for the invention may admit to other equally effective embodiments.

[0012] FIG. 1 illustrates an example networked environment in which machine learning models are trained and deployed for image matting, according to an embodiment.

[0013] FIG. 2 illustrates a pipeline for determining alpha mattes from input images and trimaps, according to an embodiment.

[0014] FIG. 3 illustrates examples of an input image and trimap, foreground and background predictions, and an alpha matte prediction, according to an embodiment.

[0015] FIG. 4 illustrates a method of training machine learning models for image matting, according to an embodiment.

[0016] FIG. 5 illustrates a method for image matting, according to an embodiment.

[0017] FIG. 6 illustrates a computing system in which an embodiment may be implemented.

DETAILED DESCRIPTION

[0018] Embodiments presented herein provide techniques for image matting. In particular, embodiments decompose the matting problem of estimating the opacity of a foreground into the targeted subproblems of estimating a background using a first trained neural network, estimating a foreground using a second neural network and the estimated background as one of the inputs into the second neural network, and estimating an alpha matte using a third neural network and the estimated background and foreground as two of the inputs into the third neural network. That is, rather than estimating foreground and background color pairs together directly for each pixel as in sampling-based matting approaches, embodiments first estimate the background image using its consistent structure, and then estimate foreground colors using the background estimates. Thereafter, the background and foreground color estimates are used, together with an input image and trimap, to estimate an alpha matte. Experience has shown that, by decomposing the matting problem into subproblems that are easier for a neural network to learn compared to traditional data-driven techniques for image matting, embodiments disclosed herein can produce better opacity estimates than traditional data-driven techniques as well as sampling-based and affinity-based matting approaches.

[0019] In the following, reference is made to embodiments of the invention. However, it should be noted that the invention is not limited to specific described embodiments. Instead, any combination of the following features and elements, whether related to different embodiments or not, is contemplated to implement and practice the invention. Furthermore, although embodiments of the invention may achieve advantages over other possible solutions and/or over the prior art, whether or not a particular advantage is achieved by a given embodiment is not limiting of the invention. Thus, the following aspects, features, embodiments and advantages are merely illustrative and are not considered elements or limitations of the appended claims except where explicitly recited in a claim(s). Likewise, reference to “the invention” shall not be construed as a generalization of any inventive subject matter disclosed herein and shall not be considered to be an element or limitation of the appended claims except where explicitly recited in a claim(s).

[0020] As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a “circuit,” “module” or “system.” Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

[0021] Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable

combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

[0022] A computer readable signal medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electromagnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.

[0023] Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, RF, etc., or any suitable combination of the foregoing.

[0024] Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the “C” programming language or similar programming languages. The program code may execute entirely on the user’s computer, partly on the user’s computer, as a stand-alone software package, partly on the user’s computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user’s computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

[0025] Aspects of the present invention are described below with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0026] These computer program instructions may also be stored in a computer readable medium that can direct a

computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

[0027] The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0028] Embodiments of the invention may be provided to end users through a cloud computing infrastructure. Cloud computing generally refers to the provision of scalable computing resources as a service over a network. More formally, cloud computing may be defined as a computing capability that provides an abstraction between the computing resource and its underlying technical architecture (e.g., servers, storage, networks), enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction. Thus, cloud computing allows a user to access virtual computing resources (e.g., storage, data, applications, and even complete virtualized computing systems) in “the cloud,” without regard for the underlying physical systems (or locations of those systems) used to provide the computing resources.

[0029] Typically, cloud computing resources are provided to a user on a pay-per-use basis, where users are charged only for the computing resources actually used (e.g. an amount of storage space consumed by a user or a number of virtualized systems instantiated by the user). A user can access any of the resources that reside in the cloud at any time, and from anywhere across the Internet. In the context of the present invention, a user may access applications (e.g., a matting application) or related data available in the cloud. For example, a matting application could execute on a computing system in the cloud to perform image matting, and store alpha mattes determined via such matting at a storage location in the cloud. Doing so allows a user to access this information from any computing system attached to a network connected to the cloud (e.g., the Internet).

[0030] FIG. 1 illustrates an example networked environment 100 in which machine learning models are trained and deployed for image matting, according to an embodiment. As shown, the networked environment 100 includes a model generator 110, a training data repository 120, an application server 130, and a data repository 140 communicatively connected via a network 150.

[0031] The model generator 110 is generally representative of a computing system, such as the system discussed below with respect to FIG. 6. As shown, the model generator 110 includes a training data generating application (“training data generator”) 112 and a model training application (“model trainer”) 114 running therein. Although shown as being distinct, in some embodiments the training data generator 112 and the model trainer 114 may be implemented as a single application. The training data generator 112 is

configured to generate training data that may be stored in the training data repository 120 (or elsewhere) and that is used, in part, by the model trainer 116 to train machine learning models for image matting. In one embodiment, training data may be generated via a data augmentation scheme that creates composite images on the fly using randomizations, as discussed in greater detail below.

[0032] The model trainer 116 is configured to train machine learning models, which may include a background sampling network, a foreground sampling network, and matting network. As detailed below, the background sampling network may take as inputs an image and associated trimap and output a predicted background; in turn, the foreground sampling network may take as input the image and trimap, as well as the background predicted by the background sampling network, and output a predicted foreground; then, the matting network may take as inputs the image and trimap, as well as the background predicted by the background sampling network and the foreground predicted by the foreground sampling network, and output a predicted alpha matte. Examples of network architectures for the background sampling network, foreground sampling network, and matting network, as well as training algorithms and loss functions to be minimized during training, are discussed in greater detail below.

[0033] The application server 130 is generally representative of another computing system, the components of which may be similar to those of the model generator 110. As shown, the application server 130 includes a matting application 132 running therein. In one embodiment, the matting application 132 is configured to sequentially utilize the trained background sampling network, the trained foreground sampling network, and the trained matting network to determine an alpha matte given an input image and associated trimap, which may be retrieved from the data repository 140 (or elsewhere), as discussed in greater detail below.

[0034] FIG. 2 illustrates a pipeline 200 for determining alpha mattes from input images and trimaps, according to an embodiment. As described, image matting involves estimating the soft transitions between a user-defined foreground and a background of an image, with the soft transitions defining the opacity of the foreground at each pixel. The goal of matting is thus to extract from a given image a foreground layer with correct opacity values and color, such that the extracted foreground layer can be used in compositing operations down the pipeline.

[0035] Formally, the color mixtures in soft transitions between foreground and background may be represented with the compositing equation for a given pixel:

$$I_i \alpha_i F_i + (1 - \alpha_i) B_i, \quad (1)$$

where $\alpha_i \in [0, 1]$ denotes the opacity of the foreground at pixel i and I , F , and B represent color values of the original image, the foreground, and the background, respectively. The matting problem is underconstrained, as the observed color of the pixel I_i in the composite image is known, while the foreground color F_i , background color B_i , and opacity α_i are to be determined. In addition to the input image I itself, a user may provide as input a trimap which indicates, for every pixel, whether the pixel is a foreground (i.e., opaque, $\alpha=1$) or background ($\alpha=0$), or of unknown opacity where alpha values need to be estimated. Embodiments aim to estimate the unknown opacities by making use of the colors

of pixels that have been classified by the user as foreground and background in the known-opacity regions.

[0036] As shown, the pipeline 200 deconstructs the solution of equation (1) by first estimating the background and foreground parameters in equation (1) using background and foreground sampling networks 230 and 250, respectively, and then estimating opacity using a matting network 270 that takes the estimated foreground and backgrounds as additional input channels, along with an input image 210 and trimap 220. This data-driven approach to color sampling employs targeted neural networks, namely the sampling networks 230 and 250, that estimate background and foreground layer colors while making use of high-level image structure and texture. Traditional sampling-based matting does not make use of the high spatial correlation between colors of neighboring pixels, whereas embodiments exploit such spatial correlations by estimating per-pixel samples while accounting for textures that exist in the foreground and background regions through the use of neural networks. Experience has shown that foreground and background samples estimated by techniques disclosed herein are higher-quality than those estimated using sampling-based matting approaches. Further, by estimating the foreground and background color parameters in equation (1) using sampling networks according to embodiments disclosed herein, the difficulty of estimating corresponding alpha values using a trained neural network is reduced compared to training a neural network to estimate alpha values given only an input image and trimap.

[0037] The intuition behind first estimating the background image 240 using its consistent structure, and then estimating the foreground 260 using the background estimates, as opposed to estimating the foreground and background color pairs directly in a single shot for each pixel as in sampling-based approaches, is as follows. As a consequence of the definition of the matting problem, the foreground and background have different characteristics. In particular, the background image may be thought of as a fully opaque layer that is partially obscured by foreground object(s). In other words, a background image with consistent structural and textural properties may be considered to be behind the foreground. By contrast, the foreground layer is spatially limited to the extent of non-opaque regions. This means that colors that participate in mixtures with the background may be expected to have similar color distributions to the fully opaque pixels, but the structural and textural characteristics may differ greatly in partial-opacity regions. Given this intuition, embodiments decompose the matting problem as described above, instead of estimating the foreground and background color pairs together directly for each pixel as in traditional sampling-based approaches.

[0038] As shown, the background sampling network 230 takes as inputs the input image 210 and trimap 220 and outputs a predicted background 240. Panel A of FIG. 3 illustrates an example input image 310 and associated trimap 320, and panel B shows an example predicted background 330 that the background sampling network 230 may output given the input image 310 and trimap 320. As shown, the input image 310 depicts a pineapple 315, and the trimap 320 is a user-generated segmentation of the image 310 that includes a region 322, shown in white, that the user is confident is foreground; a region 324, shown in black, that the user is confident is background; and a region 326, shown in gray, whose pixels are not known to be foreground or

background and could be a combination of the two (in which the foreground is partially transparent). The goal is then to determine alpha values representing opacity of foreground object(s) in grey areas such as the region 326.

[0039] In one embodiment, the background sampling network 230 is a neural network with a two-stage architecture in which both stages are fully convolutional encoder-decoder structures, although other network architectures, such as architectures suitable for inpainting or sampling more generally, may be used in alternative embodiments. For example, a two-stage architecture from image inpainting, which attempts to hallucinate a missing part of an image using the consistent structures and high-level understanding of the overall image, may be used. However, unlike the image inpainting problem, the background image that the background sampling network 230 estimates is partially observed in the input image hidden behind the partially opaque foreground with an unknown alpha matte. To keep the two-stage architecture from hallucinating plausible structures in the unknown region, the background sampling network 230 may be trained to recognize the background structures that are hidden behind the foreground and estimate samples accordingly. In one embodiment, this is achieved by providing the input image 210 as one of the inputs to the background sampling network 230, rather than only providing fully transparent regions defined by the trimap 220. In such a case, the loss function used during training of the background sampling network 230 may be defined over only the unknown-opacity region, as:

$$\mathcal{L}_B = \frac{1}{|U|} \sum_{i \in U} |\hat{B}_i - B_i|, \quad (2)$$

where \hat{B} and B denote the predicted and ground-truth background colors, and U is the image region labeled as unknown in the input trimap.

[0040] With the additional input of the input image 210 and the use of equation (2), the background sampling network 230 learns not to hallucinate regions but to instead use the partially obstructed background regions as a guide to estimate high-frequency details. This difference is crucial for sampling-based matting, as color values for the background directly affect the matte quality through its use in the compositing equation. It should also be noted that the loss function of equation (2) only includes the unknown-opacity region, as the background colors are irrelevant for the fully-opaque foreground regions in alpha estimation.

[0041] In a particular embodiment, the background sampling network 230 may be a neural network with a two-stage architecture. Both stages are fully convolutional encoder-decoder structures. The first stage may take an RGB (red, green, blue) image and a binary mask indicating the inpainting regions as input pairs and output the initial coarse inpainted RGB image. The whole RGB image may be taken as input, rather than the image with white pixels filling in holes, as is typically the case with normal image inpainting tasks. Doing so shifts the network's functionality of purely hallucinating missing colors to instead distinguishing background colors from mixed input colors through further training with whole RGB image input. In addition, for the image matting task, the binary mask may come directly from the trimap 320, and the unknown and foreground regions of the trimap 320 are the inpainting region during background

sampling. The second stage of the two-stage architecture further takes the coarse prediction from the first stage along with the binary mask as inputs and predicts final results. The second stage may include two parallel encoders: a dilational convolution branch specifically focused on hallucinating contents with layer-by-layer dilated convolution, and a contextual attention branch that tries to attend on distant background features of interest. In terms of implementation, the background sampling network **230** may use zero-padding for each convolutional layer, with no batch normalization layers, the exponential linear unit (ELU) may be used rather than the rectifier linear unit (ReLU) as the activation function, output filter values may be clipped rather than activated by tanh or sigmoid functions, and the Nearest Neighbor method may be used in the upsample layer.

[0042] Once background samples are estimated, the foreground sampling problem boils down to choosing plausible colors from the known fully-opaque foreground regions that best represent color mixtures observed in the input image **210**. Background sampling is performed prior to foreground sampling in embodiments, as the background may be more uniform than the foreground which may have, e.g., varied textures, making the likelihood higher of finding better background samples. Returning to FIG. 2, the foreground sampling network **250** receives the input image **210** and trimap **220**, as well as the predicted background **240**, as inputs. In turn, the foreground sampling network **250** processes such inputs and outputs a predicted foreground **260**. Panel B of FIG. 3 illustrates an example foreground **340** that the foreground sampling network **240** may output given the input image **310**, the trimap **320**, and the predicted background **330**.

[0043] In one embodiment, the foreground sampling network **250** is, like the background sampling network **230**, a neural network with a two-stage architecture in which both stages are fully convolutional encoder-decoder structures, although other network architectures, such as architectures suitable for inpainting or sampling more generally, may be used in alternative embodiments. In such a case, the loss function used to train the foreground sampling network **250** may be composed of two terms added together:

$$\mathcal{L}_F = \mathcal{L}_{fg} + \mathcal{L}_{comp} \quad (3)$$

[0044] The first term \mathcal{L}_{fg} in the loss function of equation (3) is the L^1 loss analogous to background sampling:

$$\mathcal{L}_{fg} = \frac{1}{|U|} \sum_{i \in U} |\hat{F}_i - F_i|, \quad (4)$$

where \hat{F} and F denote the predicted and ground-truth foreground colors. This term simply changes the binary mask to assign background and unknown area of trimap **320** to be the inpainting region and incurs loss on the differences between ground-truth and predicted foreground colors in the unknown region U . However, sampling foreground is more challenging than sampling background colors, as foreground inpainting regions are typically larger than those in background inpainting and are connected to the boundaries of images, and, unlike background colors, foreground colors are typically of high frequency. To alleviate these problems, background colors and alpha mattes may be utilized during training to help the foreground sampling network **250**. To use background colors, the foreground sampling network

250 may take background colors predicted from the fully trained background sampling network **230** as extra input channels (e.g., 7 channels rather than 4) to gain more information. In one embodiment, the extra weights in the first convolutional layer of the foreground sampling network **250** may be initialized with Xavier random variables. To use alpha mattes, an additional compositional loss may be added to the foreground prediction loss \mathcal{L}_f to guide the network to learn predicting foreground colors in the manner defined by equation (1). The second term in the loss function of equation (3) is such a compositional loss that penalizes deviations of an intermediate composite image by using equation (1) with predicted background and foreground colors, and ground-truth alpha mattes from the reference composite input image I :

$$\mathcal{L}_{comp} = \frac{1}{|U|} \sum_{i \in U} |\alpha_i \hat{F}_i + (1 - \alpha_i) \hat{B}_i - I_i|. \quad (5)$$

Here, the ground-truth alpha matte, ground-truth background, and predicted foreground are used together to compute a predicted composited RGB image $\alpha_i \hat{F}_i + (1 - \alpha_i) \hat{B}_i$, and the predicted RGB image is compared with the ground-truth RGB image in L_1 distance. It should be noted that, although ground-truth alpha mattes are used during training of the foreground sampling network **250** through the compositional loss, the alpha matte is not needed when estimating foreground samples in the final system in forward passes.

[0045] In a particular embodiment, a weighted version of the compositional loss may be employed in which an ad hoc weight is applied depending on the ground-truth a pixel-wisely to the L_1 distance between predicted and ground-truth composited images. In such a case, for each pixel i , the weight may be computed as:

$$\omega^i = \begin{cases} 3\alpha^i + \frac{1}{4}, & \text{for } \alpha^i \in [0, 0.5] \\ -3\alpha^i + \frac{13}{4}, & \text{for } \alpha^i \in [0.5, 1] \end{cases}, \quad (6)$$

[0046] The weighted compositional loss may then be computed as:

$$\mathcal{L}_{wcomp} = \frac{1}{|U|} \sum_{i \in U} \omega_i |\hat{I}_i - I_i|, \quad (7)$$

where $I = \alpha \hat{F} + (1 - \alpha) \hat{B}$. Adding a weight to the compositional loss aims to increase the influence of the compositional loss term where α is close to 0.5, indicating final image colors are almost equally blended between foreground and background colors and to decrease the influence of the compositional loss term where α is close to 0 or 1, meaning the image color is almost the same as the foreground or background color. The final loss may then be the weighted sum of the foreground prediction loss and the weighted compositional loss:

$$\mathcal{L}_F = \mathcal{L}_{fg} + \lambda_{comp} \mathcal{L}_{wcomp} \quad (8)$$

[0047] Experience has shown that the background and foreground sampling networks disclosed herein are able to generate spatially smooth color predictions that match well

with actual colors that form mixtures in images, in contrast to traditional sampling-based matting approaches that generated noisy sampling results. Further, due to the common use of the color line model to select samples, traditional approaches can incorrectly choose foreground samples that are very similar to the actual background color in transparent regions, and vice versa, resulting in low-quality alpha predictions.

[0048] Returning to FIG. 2, after the predicted foreground **260** is determined using the foreground sampling network **250**, then the matting network **270** is further used to process the input image **210**, the trimap **220**, the predicted background **240**, and the predicted foreground **260**. In turn, the matting network **270** outputs the predicted alpha matte **280** corresponding to the input image **210**. Panel B of FIG. 3 illustrates an example alpha matte **350** that the matting network **270** may output given the input image **310**, the trimap **320**, the predicted background **330**, and the predicted foreground **340**.

[0049] In one embodiment, the matting network **270** is a generative adversarial network (GAN), although other network architectures such as an encoder-decoder architecture or the Deep Mating architecture may be used in alternative embodiments. A GAN may be implemented as a system of two artificial neural networks, a generative model that is trained to map inputs to a desired data distribution and a discriminative model that is trained to evaluate the probability that a sample came from training data rather than the generative model. In a particular embodiment, an AlphaGAN architecture, which includes a generative model that is a convolutional encoder-decoder network trained both with the help of ground-truth alphas as well as the adversarial loss from a discriminative network, may be used. In such a case, the input to the encoder may include 10 channels: three (red, green, and blue) from the RGB image **210**, one from the trimap **220**, three from the background colors **240**, and three from the foreground colors **260**.

[0050] In one embodiment, the matting network **270** may be trained using a loss function that includes the alpha prediction loss \mathcal{L}_{α} and compositional loss \mathcal{L}_{comp} used to train the Deep Matting network, which together are referred to herein as \mathcal{L}_{dm} , as well as an additional loss \mathcal{L}_{grad} grad that is defined over alpha gradients. The alpha prediction loss \mathcal{L}_{α} alpha may be approximated as $\mathcal{L}_{\alpha}^i = \sqrt{(\alpha_p^i - \alpha_g^i)^2 + \epsilon^2}$, where α_p^i is the output of the prediction layer at pixel i , and the compositional loss \mathcal{L}_{comp} may be approximated as $\mathcal{L}_{comp}^i = \sqrt{(c_p^i - c_g^i)^2 + \epsilon^2}$, where c is the RGB channel, p is the image composited by a predicted alpha, and g is the image composited by ground truth alphas. However, experience has shown that the alpha prediction loss \mathcal{L}_{α} and compositional loss \mathcal{L}_{comp} do not properly promote sharpness in alpha mattes. This issue is addressed using the \mathcal{L}_{grad} grad loss that is defined over alpha gradients defined as the L^1 distance between the spatial gradient of predicted and ground-truth alpha mattes:

$$\mathcal{L}_{grad} = \frac{1}{|U|} \sum_{i \in U} |\hat{G}_i - G_i|, \quad (9)$$

where \hat{G} and G denote the gradient magnitude of predicted and ground-truth alpha mattes, respectively. It should be noted that, rather than operating on color values as the alpha

prediction loss \mathcal{L}_{α} and compositional loss \mathcal{L}_{comp} do, the \mathcal{L}_{grad} loss ensures that alpha gradients are also considered during the training process. The final matting loss that includes the alpha prediction loss \mathcal{L}_{α} , the compositional loss \mathcal{L}_{comp} , and the \mathcal{L}_{grad} loss over alpha gradients may be defined as:

$$\mathcal{L}_{matting} = \mathcal{L}_{dm} + \mathcal{L}_{grad}. \quad (10)$$

Experience has shown that the loss of equation (10) promotes sharp results that effectively capture high-frequency alpha changes. Intuitively, ensuring that gradients are preserved in the alpha matte may help produce better results in regions having smooth transitions of gradients.

[0051] FIG. 4 illustrates a method **400** of training machine learning models for image matting, according to an embodiment. As shown, the method **400** begins at step **410**, where the training data generator **112** generates training data. In one embodiment, the training data generator **112** may generate training data for training the foreground and background sampling networks using similar preprocessing as for training Deep Matting network. The training data generator **112** may further generate training data for training the matting network from a set of unique foreground images and corresponding alpha mattes, and using a data augmentation scheme that creates new composite images on the fly by many randomizations. In one embodiment, the randomizations may include extending the number of distinct ground-truth mattes by compositing two random foreground images on top of each other with a probability of 0.5. The corresponding trimaps may then be defined, again on the fly, by dilating the foreground by a random number of pixels ranging from, e.g., 1 to 19. Doing so creates greater variability in terms of ground-truth mattes and trimaps, and provides better generalizability as images with sharp opacity boundaries and large transparent regions are randomly combined to create hybrid scenes. The randomizations may further include, after the foreground images are selected and mixed, selecting a number (e.g., 8) of random background images and creating composite images, which are then fed to the matting network as a batch. Using the same foreground with different backgrounds in a batch increase the invariance of the matte estimation with respect to the background image. Selecting the background images randomly at each iteration instead of defining a pre-determined set results in the network seeing new input images throughout the training process, which also helps with generalizability of the final network. In addition, the randomizations may include applying random scaling to the input images before feeding them to the network. That is, images may be randomly resized to, e.g., 480×480 or 640×640 before cropping patches from those images. As a result, the network sees not only zoomed-in versions of the input images, but also the input images as a whole, allowing the network to better generalize in terms of image scale.

[0052] During the training of the foreground sampling network **230**, the training data generator **112** may further employ an additional processing step to enlarge the foreground set, by, e.g., enhancing each foreground image before composition by randomly changing its brightness, contrast, and saturation within the range [0.5, 1.5] according to a truncated Gaussian distribution with $\mu=1$ and $\sigma=0.2$. The hue of the foreground image may also be randomly shifted within the range [-0.2, 0.2] according to another truncated Gaussian distribution with $\mu=0$ and $\sigma=0.1$.

[0053] At step 420, the model trainer 114 trains a background sampling network using the training data generated at step 410 (using the Deep Matting preprocessing for the background sampling network). In one embodiment, the background sampling network may have a two-stage architecture, as described above, and be trained using the loss function of equation (2) that is only defined over unknown-opacity region(s) (indicated by trimaps associated with images in the training data). Any suitable training algorithm may be used, depending on the type of background sampling network chosen. In a particular embodiment, the training may start from a pre-trained model, such as the inpainting model discussed above, and use the Adam optimizer with $\beta_1=0.9$ and $\beta_2=0.999$ for back propagation, with a constant learning rate of, e.g., 10^{-5} .

[0054] At step 430, the model trainer 114 trains a foreground sampling network using the training data generated at step 410 (using the Deep Matting preprocessing for the foreground sampling network) and background predictions made by the trained background sampling network. In one embodiment, the foreground sampling network may have a two-stage architecture, and be trained using the loss function of equation (3) that includes a L^1 loss of equation (4), which is defined over unknown-opacity region(s), and the compositional loss of equation (5), which penalizes deviations of an intermediate composite image, as described above. Similar to the training of the background sampling network at step 420, any suitable training algorithm may be used to train the foreground sampling network, depending on the type of foreground sampling network chosen, and in a particular embodiment the training may start from a pre-trained model and use the Adam optimizer with a constant learning rate.

[0055] At step 440, the model trainer 114 trains a matting network using the training data generated at step 410 (for training the matting network), as well as background and foreground predictions made by the trained background and foreground sampling networks, respectively. In particular, the matting network is trained to output a single channel for the alpha matte. In one embodiment, the matting network is trained using the loss function of equation (10), which includes an Alpha prediction loss \mathcal{L}_{α} , a compositional loss \mathcal{L}_{comp} , and the \mathcal{L}_{grad} loss of equation (9) that is defined as the L^1 distance between the spatial gradient of predicted and ground-truth alpha mattes, as discussed above. Once again, any suitable training algorithm may be used to train the matting network, depending on the type of matting network chosen.

[0056] In some embodiments, the matting network may further be trained to be specialized for specific types of scenes. For example, additional training using data from a specific film franchise could be used to improve the matte quality generated for image frames of films of the same franchise. Here, it is assumed that the overall appearance of image frames from the same film franchise are similar to each other, such that training within this relatively lower dimensional space will be easier than training with completely random images. In particular, if the post-processing of films from the same film franchise were performed in a similar manner, such as using the same color grading, or if the same characters or objects appear in the films, the dimensionality of the problem may be reduced. For example, if the networks were to be trained on a series of films in which most shots are shot with similar lighting, are

color graded the same way, involve the same characters, take place in similar places, etc., then the networks may work better on a future installment of the same film franchise that involves the same characteristics. In one embodiment, the model trainer 114 may start with already trained networks, and train those networks further with new data such as data from a specific film franchise.

[0057] FIG. 5 illustrates a method 500 for image matting, according to an embodiment. As shown, the method 500 begins at step 510, where the matting application 132 receives an input image and an associated trimap. As described, the trimap may indicate foreground region(s), background region(s), and region(s) of unknown opacity in the input image, with the goal of matting being to predict the opacity of the foreground at each pixel in the unknown opacity region(s).

[0058] At step 520, the matting application 132 processes the input image and trimap using a trained background sampling network to predict a background in the input image. As described, the background sampling network takes an RGB image and trimap as inputs and is trained to fill in missing background colors, thereby predicting the background of the input image.

[0059] At step 530, the matting application 132 processes the input image, trimap, and predicted background using a trained foreground sampling network to predict a foreground in the input image. As described, the foreground sampling network is trained to take an RGB image and trimap, along with the background predicted at step 520, as inputs and to predict foregrounds colors given such inputs.

[0060] At step 540, the matting application 132 processes the input image, trimap, predicted background, and predicted foreground using a trained matting network to predict an alpha matte. As described, the matting network uses all available information, including the RGB image, the trimap, as well as the background predicted at step 520 and the foreground predicted at step 530, to predict an alpha matte. That is, the foreground and background predictions made at steps 520 and 530 are input into the matting network as image channels, in addition to the RGB and trimap image channels. In turn, the matting network outputs a single channel for the alpha matte.

[0061] FIG. 6 illustrates a computing system 600 in which an embodiment may be implemented. As shown, the computing system 600 combines the model generator 110 and the application server 130 discussed above for simplicity, but it should be understood that the model generator 110 and the application server 130 may also be implemented as distinct computing systems. Illustratively, the system 600 includes, without limitation, a central processing unit (CPU) 605, a network interface 615 connecting the system to a network 616, an interconnect 617, a memory 620, and storage 630. The system 600 may also include an I/O device interface 610 connecting I/O devices 612 (e.g., keyboard, display and mouse devices) to the system 600.

[0062] The CPU 605 retrieves and executes programming instructions stored in the memory 620. Similarly, the CPU 605 stores and retrieves application data residing in the memory 620. The interconnect 617 facilitates transmission, such as of programming instructions and application data, between the CPU 605, I/O device interface 610, storage 630, network interface 615, and memory 620. CPU 605 is included to be representative of a single CPU, multiple CPUs, a single CPU having multiple processing cores, one

or more graphics processing units (GPUs), a combination of the above, and other types of processor(s). And the memory 620 is generally included to be representative of a random access memory. The storage 630 may be a disk drive storage device. Although shown as a single unit, the storage 630 may be a combination of fixed and/or removable storage devices, such as magnetic disk drives, flash drives, removable memory cards or optical storage, network attached storage (NAS), or a storage area-network (SAN). Further, the system 600 is included to be representative of a physical computing system as well as virtual machine instance(s) hosted on underlying physical computing system(s). Further still, although shown as a single computing system, one of ordinary skill in the art will recognize that the components of the system 600 shown in FIG. 6 may be distributed across multiple computing systems connected by a data communications network.

[0063] As shown, the memory 620 includes an operating system 621, the training data generator 112, the model trainer 114, and the matting application 132. The operating system 621 may be, e.g., Linux® or Microsoft Windows®. As described, the training data generator 112 is configured to generate training data; the model trainer 114 is configured to sequentially train a background sampling network, a foreground sampling network, and a matting network; and the matting application 132 is configured to process an input image and associated trimap using the trained background and foreground sampling networks and the trained matting network in order to determine an alpha matte. In one embodiment, the training data generator 112 generates training data, after which the model trainer 114 trains a background sampling network using the generated training data; trains a foreground sampling network using the generated training data and background predictions made by the trained background sampling network; and trains a matting network using the generated training, as well as background and foreground predictions made by the trained background and foreground sampling networks, respectively, according to the method 400 discussed above with respect to FIG. 4. After the foreground and background sampling networks and the matting network have been trained, the matting application 132 may receive an input image and an associated trimap; process the input image and trimap using the trained background sampling network to predict a background in the input image; process the input image, trimap, and predicted background using the trained foreground sampling network to predict a foreground in the input image; and process the input image, trimap, predicted background, and predicted foreground using the trained matting network to predict an alpha matte, according to the method 500 discussed above with respect to FIG. 5.

[0064] Advantageously, techniques disclosed herein permit image matting with learning-based sampling. Sampling networks are disclosed for estimating background and foreground layer colors while making use of high-level image structure and texture, in contrast to the limited set of selected samples used by traditional sampling-based matting approaches. Experience has shown that by inputting foreground and background samples estimated using such sampling networks into a matting neural network, along with an input image and trimap, techniques disclosed herein can achieve improvements in opacity estimation over traditional hand-crafted sampling approaches and affinity-based matting, as well as previous data-driven techniques that directly

solve the compositing equation for alpha values. The matting network itself may also be trained using a loss that is defined over alpha gradients, which helps promote proper sharpness in alpha mattes (which may effectively capture high-frequency alpha changes) produced by the matting network, in contrast to the traditional alpha prediction loss and compositional loss that have been used. Techniques disclosed herein may fix various issues in traditional sampling-based matting and affinity-based matting approaches and data-driven techniques, such as the inability of traditional sampling approaches to produce satisfactory results in many cases, the problems encountered with affinity-based matting when the background or foreground colors in an unknown region does not clearly appear in the known regions of a trimap, and the smoothness issues and erroneous low-alpha values in high-transparency images produced by traditional data-driven techniques. Further, the accuracy of alpha mattes and foreground colors produced by techniques disclosed herein may enable practical compositing applications, such as foreground extraction and background replacement. Additional advantages include the possibility of implementing techniques disclosed herein without compromising matte quality, with embodiments potentially being ported to various platforms (including platforms that can benefit from improved processing hardware). Interactive computation rates are also possible, as the processing amounts to a forward pass step of neural networks. In addition, techniques disclosed herein may be used to train neural networks that are specialized for specific types of scenes.

[0065] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order or out of order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

[0066] While the foregoing is directed to embodiments of the present invention, other and further embodiments of the invention may be devised without departing from the basic scope thereof, and the scope thereof is determined by the claims that follow.

What is claimed is:

1. A computer-implemented method for image matting, the method comprising:

processing a received image and associated trimap using, at least in part, a first machine learning model which outputs a predicted background;

processing the received image, the associated trimap, and the predicted background using, at least in part, a second machine learning model which outputs a predicted foreground; and

processing the received image, the associated trimap, the predicted background, and the predicted foreground using, at least in part, a third machine learning model which outputs an alpha matte.

2. The method of claim 1, wherein:

the first machine learning model is trained using at least a portion of training data which includes images and corresponding trimaps;

the second machine learning model is trained using at least a portion of the training data and backgrounds predicted by the first machine learning model subsequent to the training of the first machine learning model; and

the third machine learning model is trained using at least a portion of the training data, backgrounds predicted by the first machine learning model, and foregrounds predicted by the second machine learning model subsequent to the training of the second machine learning model.

3. The method of claim 2, wherein the training of the third machine learning model attempts to minimize a loss function which includes a loss defined over alpha gradients.

4. The method of claim 3, wherein the loss defined over alpha gradients is defined as a L^1 distance between a spatial gradient of predicted and ground-truth alpha mattes.

5. The method of claim 2, wherein the training of the first machine learning model attempts to minimize a loss function which is defined over one or more unknown-opacity regions indicated by the trimaps corresponding to the images included in the training data.

6. The method of claim 2, wherein the training of the second machine learning model attempts to minimize a loss function which includes a L^1 loss defined over one or more unknown-opacity regions and a compositional loss which penalizes deviations of an intermediate composite image.

7. The method of claim 2, further comprising:

generating the training data which includes the images and the corresponding trimaps,

wherein the generating of the images in the training data includes at least one of compositing pairs of random foreground images on top of each other with a pre-defined probability, selecting random background images and creating composite images using the selected background images and the composited pairs of random foreground images, applying random scaling, or randomly changing at least one of foreground image brightness, contrast, saturation, or hue.

8. The method of claim 1, wherein each of the first and second machine learning models has a two-stage architecture in which both stages are fully convolutional encoder-decoder structures.

9. The method of claim 1, wherein the third machine learning model has one of a generative adversarial network (GAN), a Deep Matting, or an encoder-decoder architecture.

10. The method of claim 1, wherein the associated trimap indicates at least one foreground region, at least one background region, and at least one region of unknown opacity in the received image.

11. A non-transitory computer-readable storage medium storing a program, which, when executed by a processor performs operations for image matting, the operations comprising:

processing a received image and associated trimap using, at least in part, a first machine learning model which outputs a predicted background;

processing the received image, the associated trimap, and the predicted background using, at least in part, a second machine learning model which outputs a predicted foreground; and

processing the received image, the associated trimap, the predicted background, and the predicted foreground using, at least in part, a third machine learning model which outputs an alpha matte.

12. The computer-readable storage medium of claim 11, wherein:

the first machine learning model is trained using at least a portion of training data which includes images and corresponding trimaps;

the second machine learning model is trained using at least a portion of the training data and backgrounds predicted by the first machine learning model subsequent to the training of the first machine learning model; and

the third machine learning model is trained using at least a portion of the training data, backgrounds predicted by the first machine learning model, and foregrounds predicted by the second machine learning model subsequent to the training of the second machine learning model.

13. The computer-readable storage medium of claim 12, wherein the training of the third machine learning model attempts to minimize a loss function which includes a loss defined over alpha gradients.

14. The computer-readable storage medium of claim 11, wherein each of the first and second machine learning models has a two-stage architecture in which both stages are fully convolutional encoder-decoder structures.

15. The computer-readable storage medium of claim 11, wherein the third machine learning model has one of a generative adversarial network (GAN), a Deep Matting, or an encoder-decoder architecture.

16. A computer-implemented method of training machine learning models for image matting, the method comprising:

training a first machine learning model using at least a portion of training data which includes images and corresponding trimaps;

training a second machine learning model using at least a portion of the training data and backgrounds predicted by the first machine learning model subsequent to the training of the first machine learning model; and

training a third machine learning model using at least a portion of the training data, backgrounds predicted by the first machine learning model, and foregrounds predicted by the second machine learning model subsequent to the training of the second machine learning model.

17. The method of claim 16, wherein:

the training of the third machine learning model attempts to minimize a loss function which includes a loss defined over alpha gradients; and

the loss defined over alpha gradients is defined as a L^1 distance between a spatial gradient of predicted and ground-truth alpha mattes.

18. The method of claim **16**, wherein the training of the first machine learning model attempts to minimize a loss function which is defined over one or more unknown-opacity regions indicated by the trimaps corresponding to the images included in the training data.

19. The method of claim **16**, wherein the training of the second machine learning model attempts to minimize a loss function which includes a L^1 loss defined over one or more unknown-opacity regions and a compositional loss which penalizes deviations of an intermediate composite image.

20. The method of claim **16**, further comprising:
generating the training data which includes the images and the corresponding trimaps,

wherein the generating of the images in the training data includes at least one of compositing pairs of random foreground images on top of each other with a predefined probability, selecting random background images and creating composite images using the selected background images and the composited pairs of random foreground images, applying random scaling, or randomly changing at least one of foreground image brightness, contrast, saturation, or hue.

21. A method of generating a training data set for training an image matting machine learning model, the method comprising:

receiving a plurality of foreground images and alpha mattes corresponding to the plurality of foreground images;

generating a plurality of composited foreground images by compositing randomly selected foreground images in the plurality of foreground images on top of each other with a predefined probability;

generating a respective trimap for each composited foreground image of the plurality of composited foreground images by dilating a foreground in the composited foreground image by a random number of pixels;

receiving a plurality of background images; and

generating composited images using random selected background images from the plurality of background images and the composited foreground images,

wherein the image matting machine learning model is trained based, at least in part, on the composited images and corresponding trimaps from the generated trimaps.

22. The method of claim **21**, further comprising:

randomly changing at least one of foreground image brightness, contrast, saturation, or hue of the plurality of foreground images to enlarge the training data set.

23. The method of claim **21**, further comprising:

applying random scaling to the generated composited images,

wherein cropped patches from the randomly-scaled composited images and the corresponding trimaps from the generated trimaps are used as the training data set during training of the image matting machine learning model.

* * * * *